

Functionals on Triangulations of Delaunay Sets *

Nikolay P. Dolbilin[†], Herbert Edelsbrunner[‡], Alexey Glazyrin[§] and Oleg R. Musin[§]

Abstract

We study densities of functionals over uniformly bounded triangulations of a Delaunay set of vertices, and prove that the minimum is attained for the Delaunay triangulation if this is the case for finite sets.

Keywords. Delaunay sets, triangulations, Delaunay triangulations, uniformly bounded triangulations, functionals, densities.

1 Introduction

A *Delaunay set* $X \subseteq \mathbb{R}^d$ has positive numbers $r < R$ such that every open ball of radius r contains at most one point, and every closed ball of radius R contains at least one point of X . Such sets were introduced as (r, R) -systems by Boris N. Delaunay in 1924. By a *triangulation* of X , we mean a simplicial complex, T , whose vertex set is X and whose underlying space is \mathbb{R}^d . This triangulation is *uniformly bounded* if there is a real number $q = q(T)$ such that the circumsphere of every d -simplex in T has radius smaller than or equal to q . A particular triangulation is the *Delaunay triangulation*, denoted as $\text{Del } X$, whose d -simplices satisfy the additional condition that all other vertices lie outside their circumspheres. It exists if X is generic, as will be explained shortly. The Delaunay triangulation of a Delaunay set is necessarily uniformly bounded. We also consider Delaunay triangulations of *finite* sets of points, for which the underlying space of the simplicial complex is the convex hull of the points.

Writing \mathcal{S}_d for the set of all d -simplices in \mathbb{R}^d , we consider functionals $F : \mathcal{S}_d \rightarrow \mathbb{R}$ for which there are constants $e = e(r, q, d)$ and $E = E(r, q, d)$ such that $e \leq F(\sigma) \leq E$ for every d -simplex σ whose edges are longer than or equal

to $2r$ and whose circumsphere has a radius smaller than or equal to q . Writing \mathcal{E} for this class of functionals, we define subclasses $\mathcal{G} \subseteq \mathcal{F} \subseteq \mathcal{E}$ by requiring additional conditions. Briefly, F belongs to \mathcal{F} if the sum of values over the d -simplices of the Delaunay triangulation of $d + 2$ points in \mathbb{R}^d is smaller than or equal to the sum over the d -simplices in the other triangulation, and F belongs to \mathcal{G} if it satisfies a similar condition for all finite sets of points. For a triangulation, we define the *density* of the functional by taking sums and lower limits over a growing sequence of balls:

$$f(T) = \liminf_{\alpha \rightarrow \infty} \frac{1}{\text{Vol}(\mathbb{B}_\alpha)} \sum_{\mathbb{B}_\alpha \supseteq \sigma \in T} F(\sigma), \quad (1)$$

where \mathbb{B}_α is the closed ball with radius α and center at the origin of \mathbb{R}^d , and $\text{Vol}(\mathbb{B}_\alpha)$ is its volume. With these definitions, we can give our main result:

- in \mathbb{R}^d , $F \in \mathcal{G}$ implies that the Delaunay triangulation minimizes the density of F among all uniformly bounded triangulations of a Delaunay set, and in \mathbb{R}^2 , $F \in \mathcal{F}$ suffices to reach the same conclusion.

There are many concrete functionals studied in the literature to which our result applies. Here, we just mention two:

- the functional that maps every triangle in \mathbb{R}^2 to the radius of its circumcircle; see [8],
- the functional that maps every d -simplex to the sum of squares of its edge lengths times the volume; see [11].

The remainder of this paper presents the detailed results in two sections.

2 Background

In this section, we introduce the background on Delaunay sets, their uniformly bounded triangulations, and functionals on such triangulations.

*This research is partially supported by the Russian Government under the Mega Project 11.G34.31.0053, RFBR grant 11-01-00735, DMS 1101688, and the European Science Foundation (ESF) under the Research Network Programme.

[†]Steklov Mathematics Institute, Moscow, Russian Federation.

[‡]IST Austria (Institute of Science and Technology Austria), Klosterneuburg, Austria, Departments of Computer Science and of Mathematics, Duke University, Durham, North Carolina, and Geomagic, Research Triangle Park, North Carolina.

[§]Mathematics Department, University of Texas at Brownsville, Texas, USA.

2.1 Delaunay Sets

We recall from Section 1 that $X \subseteq \mathbb{R}^d$ is a *Delaunay set* if there are positive constants $r < R$ such that (I) every open ball of radius r contains at most one point of X , and (II) every closed ball of radius R contains at least one point of X . Hence, X has no tight cluster and leaves no large hole.

Counting points. Condition (I) implies that every bounded subset of \mathbb{R}^d contains only finitely many points of X . Indeed, the subset can be covered by finitely many open balls of radius r , and each such ball contains at most one point. Condition (II) implies that every cone with non-zero volume contains infinitely many points of X . Indeed, the cone contains an infinite string of disjoint closed balls of radius R , and each such ball contains at least one point of X . We quantify the first observation by giving concrete estimates. Let $\mathbb{B}_\alpha(z)$ be the closed ball with radius α and center z , and call the difference between two concentric balls an *annulus*.

POINT COUNT LEMMA. Let X be a Delaunay set with parameters $r < R$ in \mathbb{R}^d .

- (i) There are constants $p = p(R, d)$ and $P = P(r, d)$ such that the number of points of X in $\mathbb{B}_\alpha(z)$ is between $p\alpha^d$ and $P\alpha^d$.
- (ii) There is a constant $P' = P'(r, d)$ such that the number of points of X in $\mathbb{B}_{\alpha+1}(z) - \mathbb{B}_\alpha(z)$ is at most $P'\alpha^{d-1}$.

PROOF. To prove (i), we note that $\mathbb{B}_\alpha(z)$ can be covered by some constant times $(\alpha/r)^d$ balls of radius r , and that we can pack some other constant times $(\alpha/R)^d$ balls of radius R in it. The lower and upper bounds follow.

To prove (ii), we cover the annulus with a constant times α^{d-1}/r^d balls of radius r . The upper bound follows.

It should be clear that the bound in (ii) also holds for annuli of constant width, but not for annuli whose width is a positive fraction of the radius.

Delaunay triangulations. Following the original idea of Boris N. Delaunay, we consider d -simplices with vertices from X such that the open ball bounded by the $(d-1)$ -dimensional circumsphere contains no points of X . We call such d -simplices *empty*. Here, it is convenient to assume that X is *generic* in the sense that no $d+2$ points in X lie on a common $(d-1)$ -sphere. Under this assumption, the empty d -simplices fit together without gap and overlap. Now consider two not necessarily empty but non-overlapping d -simplices that share a $(d-1)$ -simplex, which is a face of both. Assuming the two d -simplices belong to a triangulation, we call this face *locally Delaunay* if the $(d+1)$ -st vertex of the second d -simplex lies outside the circumsphere of the first d -simplex. Note that the condition is symmetric because the two circumspheres intersect in the $(d-2)$ -sphere that passes through the vertices of the face, and either both $(d+1)$ -st vertices lie outside or both

lie inside the respective other circumsphere. Delaunay considered both conditions and proved that they are equivalent [2].

DELAUNAY TRIANGULATION THEOREM. Let X be a generic Delaunay set in \mathbb{R}^d .

- (i) The collection of empty d -simplices together with their faces form a triangulation of X , commonly known as the *Delaunay triangulation*, $\text{Del } X$.
- (ii) If all $(d-1)$ -simplices of a triangulation T of X are locally Delaunay, then $T = \text{Del } X$.

The equivalence between the local and the global conditions expressed in (ii) also holds for finite sets X . In the plane, it means that a triangulation of a generic set X is Delaunay iff for each edge the sum of opposite angles in the two incident triangles is less than π .

2.2 Uniformly Bounded Triangulations

Let X be a generic Delaunay set in \mathbb{R}^d , and let T be a triangulation of X . We recall that this means that T is a simplicial complex with vertex set X whose underlying space is \mathbb{R}^d . Recall also that T is *uniformly bounded* if there is a real number $q = q(T)$ such that the radius of the circumsphere of every d -simplex in T is smaller than or equal to q . It follows that no edge of T is longer than $2q$. Note that the Delaunay triangulation of X is uniformly bounded with $q = R$.

Not every triangulation is uniformly bounded. We begin by showing that every Delaunay set has triangulations that are not uniformly bounded. Given X , we construct such a triangulation in three steps.

1. For every point $x \in X$ and every $L > 0$, we can find many points $y \in X$ such that the edge xy is longer than L and does not pass through any other points of X . Indeed, there is such an edge near every direction out of x . To see this, we consider the set of points in X that lie within the closed ball of radius L around x . There are only finitely many such points, which implies that within each cone with non-zero volume and apex x , we can find a subcone, again with non-zero volume and apex x , that does not contain any of the points of X inside the ball. However, as argued in Section 2.1, the cone contains infinitely many points of X , so they must all be at distance larger than L from x . Among these, let y be the point closest to x .

Using the knowledge about long edges, we construct a triangulation inductively, one phase at a time. After the k -th phase, we will have a triangulation T_k of a finite subset of X that includes all points at distance k or less from the origin. In addition, we will make sure that T_k contains at least one edge longer than k , and that every point of X that belongs to the underlying space of T_k is a vertex of T_k . We start with T_0 consisting of a single edge connecting the two points of X that are closest to the origin of \mathbb{R}^d .

2. In the $(k+1)$ -st phase, we let x be a vertex in the boundary of T_k . Let y be another point of X such that the edge xy is longer than $k+1$ and does not intersect the simplices in T_k other than at x . Since T_k is a triangulation, its underlying space is convex, and its boundary is triangulated. Let σ be an i -simplex in the boundary of T_k that is visible from y . We extend T_k by adding the $(i+1)$ -simplex formed by y and the vertices of the i -simplex. Doing this for y and all visible simplices in the boundary of T_k , we obtain a simplicial complex T'_k by *starring* from y . Similarly, we add a point z at distance $k+1$ or less from the origin that lies outside the underlying space by starring to the triangulation. Repeating this operation for all such points z , we eventually get a simplicial complex T''_k .
3. While T''_k is a valid simplicial complex in \mathbb{R}^d , some of the new simplices may contain points of X in their interiors. By construction, xy is not among these simplices. Let $w \in X$ be such a point, and $\sigma \in T''_k$ the simplex of lowest dimension, j , that contains w in its interior. Note that $j \geq 1$. We fix the situation by decomposing σ into $j+1$ j -simplices, each the convex hull of w and j vertices of σ . Similarly, we decompose each simplex that contains σ as a face into $j+1$ simplices of the same dimension by starring from w . Repeating this procedure for all such points w , we eventually get a triangulation T_{k+1} such that all points in X that belong to the underlying space of T_{k+1} are in fact vertices of T_{k+1} .

Observe that the edge xy added to T_k in Step 2 remains undivided until the end. This implies that T_{k+1} indeed contains an edge longer than $k+1$, as required. It follows that the triangulation thus constructed by transfinite induction is not uniformly bounded.

Measuring volume. The remainder of this section states and proves properties of uniformly bounded triangulations. We begin with the volume of their simplices.

VOLUME LEMMA. Let X be a Delaunay set with parameters $r < R$ in \mathbb{R}^d , and let T be a uniformly bounded triangulation with parameter q of X .

- (i) In \mathbb{R}^2 , there is a positive constant $v = v(r, q)$ such that $v \leq \text{Area}(\sigma)$ for every triangle σ in T .
- (ii) In \mathbb{R}^d , there is a constant $V = V(q, d)$ such that $\text{Vol}(\sigma) \leq V$ for every d -simplex σ in T .

PROOF. We prove (i) by expressing the area of a triangle in terms of the three edge lengths and the radius of the circumcircle: $\text{Area}(\sigma) = \frac{abc}{4r}$. The edges cannot be shorter than $2r$, and the radius cannot be larger than q , which implies $v(r, q) = 2r^3/q \leq \text{Area}(\sigma)$.

To prove (ii), we note that every d -simplex is contained in the ball bounded by its circumsphere. Since the radius is at most q , this ball is smaller than $V(q, d) = (2q)^d > \text{Vol}(\sigma)$.

If we remove the requirement of uniform boundedness, then the proof of the Volume Lemma breaks down. It is not clear whether the upper bound fails. In this context, we mention a related question asked by L. Danzer and independently by M. Boshernitzen: “*is it true that for every planar Delaunay set there exists a triangle with arbitrarily large area that contains no points in its interior?*” This question is still open.

Next, we describe a Delaunay set in \mathbb{R}^3 that has tetrahedra of arbitrarily small volume in the Delaunay triangulation. It shows that the limitation of the lower bound in (i) to two dimensions is necessary. Consider the standard cubic lattice, \mathbb{Z}^3 . Let $\delta_\ell = \frac{1}{2+|\ell|}$, for every $\ell \in \mathbb{Z}$, and move every point $(i, j, k) \in \mathbb{Z}^3$ to $(i, j, k + (-1)^{i+j}\delta_k)$, denoting the new point set by X . To study the volume of the tetrahedra in the Delaunay triangulation, we consider a single integer cube, for which we get a tetrahedron of volume about $\frac{1}{3}$ in the middle, four tetrahedra of volume about $\frac{1}{6}$ across each face, and two flat tetrahedra at the top and the bottom; see Figure 1. These




Figure 1: A distorted cube decomposed into seven tetrahedra. The arrows indicating the distortion are exaggerated for better visibility.

volume estimates assume arbitrarily small values of δ_ℓ . If the third coordinates of the original vertices are k and $k+1$, then the volume of the top tetrahedron is $\frac{2}{3}\delta_{k+1}$, and that of the bottom tetrahedron is $\frac{2}{3}\delta_k$. Since among the δ_ℓ there are arbitrarily small numbers, there are tetrahedra in $\text{Del } X$ whose volume is arbitrarily close to 0. It is not difficult to extend this example to four and higher dimensions.

Counting simplices. Recall the Volume Lemma, which states that every triangle in a uniformly bounded triangulation of a Delaunay set in \mathbb{R}^2 has an area that exceeds a positive constant. Since a disk of radius α has area $\alpha^2\pi$, this implies that the number of triangles contained in this disk is at most some constant times α^2 . A similar result holds in three and higher dimensions, but the lack of a lower bound on the volume of a d -simplex requires a different argument, which we present as the proof of the following bounds.

SIMPLEX COUNT LEMMA. Let X be a Delaunay set with parameters $r < R$ in \mathbb{R}^d , and let T be a uniformly bounded triangulation with parameter q of X .

- (i) The number of d -simplices sharing a common vertex is bounded from above by a constant $S'' = S''(r, q, d)$.

- (ii) There are positive constants $s = s(R, q, d)$ and $S = S(r, q, d)$ such that the number of simplices contained in a ball of radius $\alpha > 4q$ is between $s\alpha^d$ and $S\alpha^d$.
- (iii) There is a constant $S' = S'(r, q, d)$ such that the number of d -simplices contained in a ball of radius $\alpha + 1$ but not in the concentric ball of radius α is at most $S'\alpha^{d-1}$.

PROOF. To prove (i), we let x be the shared vertex, and we note that all incident d -simplices are contained in the ball of radius $2q$ centered at x . By the Point Count Lemma, the number of points in this ball is bounded from above by $P(r, d) \cdot (2q)^d$. We have at most one d -simplex for every combination of d of these points, which gives $S''(r, q, d)$.

The upper bound in (ii) is now easy: by the Point Count Lemma, the number of points inside the ball of radius α is at most $P(r, d) \cdot \alpha^d$. Multiplying $P(r, d)$ with $S''(r, q, d)$ gives $S(r, q, d)$. To get the lower bound, we restrict ourselves to the ball of radius $\alpha - 2q$. By the Point Count Lemma, the number of points in this smaller ball is at least $p(R, q) \cdot (\alpha - 2q)^d$. Every d -simplex incident to one of these points is contained in the ball of radius α . Each point belongs to at least $d + 1$ d -simplices, which implies that the lower bound on the number of points also applies to the d -simplices. Finally, $(\alpha - 2q)^d$ is at least $\alpha^d / 2^d$.

To prove (iii), we use the upper bound of $P'(r, d) \cdot \alpha^{d-1}$ on the number of points in the annulus. Each d -simplex we count is incident to at least one of these points. Multiplying $P'(r, d)$ with $S''(r, q, d)$ gives $S'(r, q, d)$.

2.3 Functionals

Recall that \mathcal{S}_d denotes the set of simplices, including degenerate ones. We are interested in functionals that have constant upper and lower bounds for the simplices that arise in uniformly bounded triangulations of Delaunay sets. For other degenerate simplices we also allow infinity as a value.

DEFINITION. Let \mathcal{E} be the class of functionals $F : \mathcal{S}_d \rightarrow \mathbb{R}$ for which there are constants $e = e(r, q, d)$ and $E = E(r, q, d)$ such that $e \leq F(\sigma) \leq E$ for all d -simplices σ with edges of length at least $2r$ and radius of the circumsphere at most q .

In this section, we extend the functionals from simplices to triangulations, and we introduce subclasses that favor Delaunay triangulations for finite sets of points.

Densities. As already mentioned in Section 1, we define the density of a functional on a triangulation by taking the lower limit over a growing ball, of the sum of values over all d -simplices in the ball divided by the volume of the ball:

$$f(T) = \liminf_{\alpha \rightarrow \infty} \frac{1}{\text{Vol}(\mathbb{B}_\alpha)} \sum_{\mathbb{B}_\alpha \supseteq \sigma \in T} F(\sigma). \quad (2)$$

There are other possibilities, such as taking the upper limit, or taking the average over the simplices. Our results extend to both modifications of the definition. One of Delaunay's motivations for defining (r, R) -systems was to generalize lattices in \mathbb{R}^d to a larger class of sets. For the Delaunay triangulation of any lattice $\Lambda \subseteq \mathbb{R}^d$, the limit of the expression in (2), in which we substitute $\text{Del } \Lambda$ for T , is well defined. Unfortunately, this is not generally the case for Delaunay triangulations of Delaunay sets, which is the reason for taking the lower limit. Since this might not be entirely obvious, we will prove shortly that for a broad class of functionals in \mathcal{E} , the limit does not generally exist. Before that, we prove some positive results, namely that the density of every functional is bounded and independent of the choice of origin. Specifically, we define

$$f_z(T) = \liminf_{\alpha \rightarrow \infty} \frac{1}{\text{Vol}(\mathbb{B}_\alpha(z))} \sum_{\mathbb{B}_\alpha(z) \supseteq \sigma \in T} F(\sigma) \quad (3)$$

for every point $z \in \mathbb{R}^d$, and we prove that all choices of z give the same result.

PROPERTIES. Let F be a functional in \mathcal{E} .

- (i) There is a constant $C = C(r, q, d)$ such that $f(T) \leq C$ for every uniformly bounded triangulation of a Delaunay set in \mathbb{R}^d .
- (ii) $f(T) = f_z(T)$ for every $z \in \mathbb{R}^d$.

PROOF. To prove (i), we recall the Simplex Count Lemma, which implies that the number of d -simplices contained in \mathbb{B}_α is bounded from above by $S(r, q, d)$. Multiplying with $E(r, q, d)$ gives $C(r, q, d)$.

To prove (ii), we let $f(T, \alpha)$ and $f_z(T, \alpha)$ be the expressions in (2) and (3) without taking the lower limit, so that $f(T) = \liminf_{\alpha \rightarrow \infty} f(T, \alpha)$, and similarly for $f_z(T)$ and $f_z(T, \alpha)$. It suffices to prove

$$\lim_{\alpha \rightarrow \infty} [f(T, \alpha) - f_z(T, \alpha)] = 0, \quad (4)$$

which we do in two steps, namely by proving

$$\lim_{\alpha \rightarrow \infty} [f(T, \alpha + L) - f(T, \alpha)] = 0, \quad (5)$$

$$\lim_{\alpha \rightarrow \infty} [f(T, \alpha + L) - f_z(T, \alpha)] = 0, \quad (6)$$

where $L = \|z\|$ is the distance between z and the origin. To prove (5), we write V_d for the volume of the unit ball in \mathbb{R}^d , and we note that $\text{Vol}(\mathbb{B}_\alpha) = V_d \alpha^d$. Furthermore, we write Σ_B and Σ_A for the sums over the d -simplices contained in the smaller ball and the extra d -simplices contained in the larger ball:

$$\Sigma_B = \sum_{\mathbb{B}_\alpha \supseteq \sigma \in T} F(\sigma), \quad (7)$$

$$\Sigma_A + \Sigma_B = \sum_{\mathbb{B}_{\alpha+L} \supseteq \sigma \in T} F(\sigma). \quad (8)$$

By the Simplex Count Lemma, Σ_B is at most some positive constant times α^d , while Σ_A is at most some constant times α^{d-1} . Hence,

$$\Delta = f(T, \alpha + L) - f(T, \alpha) \quad (9)$$

$$= \frac{\Sigma_B + \Sigma_A}{V_d(\alpha + L)^d} - \frac{\Sigma_B}{V_d\alpha^d} \quad (10)$$

$$= \frac{\Sigma_A}{V_d(\alpha + L)^d} - \frac{(\alpha + L)^d \Sigma_B - \alpha^d \Sigma_B}{V_d(\alpha + L)^d \alpha^d}. \quad (11)$$

The first term in (11) goes to zero because Σ_A grows slower than α^d , and the second term goes to zero because $(\alpha + L)^d - \alpha^d$ grows slower than α^d . The argument for (6) is similar. Indeed, all we need is to notice that the set of d -simplices contained in $\mathbb{B}_{\alpha+L}$ but not contained in $\mathbb{B}_\alpha(z)$ is a subset of those contained in $\mathbb{B}_{\alpha+L}$ but not contained in $\mathbb{B}_{\alpha-L}$. By the Simplex Count Lemma, the number of simplices thus defined is bounded from above by a constant times α^{d-1} , so that the argument goes through as before.

Non-existence of limits. We now show that taking the lower limit in the definition of density is necessary because the limit does not generally exist. Indeed, functionals for which the limit exists, even just for all Delaunay triangulations of Delaunay sets, are the exception. This is true in particular for the functionals that are invariant under isometries, which include all examples we discuss in this paper.

We begin by exhibiting a construction in \mathbb{R}^2 that acts as a stepping stone in our argument. Let σ and τ be two *compatible* triangles, by which we mean that they share an edge, the two angles opposite that edge add up to less than π , and the remaining four angles are all acute. The condition implies that at least one of the triangles is acute, and we assume σ is. Using a linear sequence of congruent copies of σ , we form a strip T_σ , which we call *wide*, and using copies of τ , we form a strip T_τ , which we call *narrow*; see Figure 2. Gluing strips




Figure 2: A Delaunay triangulation made of a block of three wide strips in the center and two blocks alternating between narrow and wide strips on both sides.

together so that they match up at boundary edges, we get a Delaunay triangulation, provided no two narrow strips are glued to each other. Let m_1, m_2, \dots be an infinite sequence

of odd integers. We construct a Delaunay triangulation D inductively, starting with a block of m_1 wide strips. On each side, we add a block of m_2 strips alternating between narrow and wide, then a block of m_3 wide strips, then a block of m_4 strips again alternating between narrow and wide, and so on. For each $i \geq 1$, let $2\alpha_i$ be the total width of the first $2i - 1$ blocks. Making sure that the origin lies on the center line of the first block, B_{α_i} is the largest disk centered at the origin that is still contained in the union of the first $2i - 1$ blocks. We consider the sequence

$$f_i = \frac{1}{\alpha_i^2 \pi} \sum_{\mathbb{B}_{\alpha_i} \supseteq \sigma \in D} F(\sigma). \quad (12)$$

Assuming that the limit in (2) exists, the sequence of f_i must converge for every sequence of m_i . This is indeed the case if we measure area, because the number of triangles that intersect \mathbb{B}_α but are not contained in it is bounded from above by a constant times α . Hence, $\lim_{\alpha \rightarrow \infty} \frac{1}{\alpha^2 \pi} \sum A(\sigma) = 1$, where we abbreviate $A(\sigma) = \text{Area}(\sigma)$ and take the sum over all triangles contained in \mathbb{B}_α , as usual. This motivates us to consider the ratios of the terms in the two sequences. Assuming \mathbb{B}_{α_i} contains k_i congruent copies of σ and ℓ_i congruent copies of τ , this gives

$$g_i = \frac{\sum_{\mathbb{B}_{\alpha_i} \supseteq \sigma \in D} F(\sigma)}{\sum_{\mathbb{B}_{\alpha_i} \supseteq \sigma \in D} A(\sigma)} = \frac{k_i F(\sigma) + \ell_i F(\tau)}{k_i A(\sigma) + \ell_i A(\tau)}. \quad (13)$$

Define $Q_\sigma = \frac{F(\sigma)}{A(\sigma)}$, $Q_\tau = \frac{F(\tau)}{A(\tau)}$, and $Q = \frac{F(\sigma) + F(\tau)}{A(\sigma) + A(\tau)}$. Assuming $Q_\sigma \neq Q_\tau$, we have $Q_\sigma \neq Q$ and define $\Delta = |Q_\sigma - Q|$. By definition, $g_1 = Q_\sigma$. We choose m_2 large enough so that $|g_2 - Q| < \frac{\Delta}{3}$, which is possible because ℓ_2/k_2 goes to 1 as m_2 goes to infinity. Then we choose m_3 large enough so that $|g_3 - Q_\sigma| < \frac{\Delta}{3}$, and so on, alternating between being close to Q_σ and Q . We thus arrive at a contradiction because there is a gap of size $\frac{\Delta}{3}$ between the terms with odd and even indices. In other words, we need $\frac{F(\sigma)}{A(\sigma)} = \frac{F(\tau)}{A(\tau)}$ for the limit to exist.

We finally show that the non-existence of the limit is not an artifact of the particular Delaunay set we used in the construction of D . Let σ' and τ' be arbitrary triangles with longest edges of lengths a and c and opposite angles 2φ and 2ψ . Setting $L > \max\{\frac{a}{2\cos\varphi}, \frac{c}{2\cos\psi}\}$, we construct triangles σ with edges of length L, L, a and τ with edges of length L, L, c . It is easy to verify that σ' and σ are compatible, and so are σ and τ , and τ and τ' . Repeating the construction with the strips and blocks three times, we see that if F is invariant under isometries of \mathbb{R}^2 and the limit exists for the Delaunay triangulations of all Delaunay sets, then $\frac{F(\sigma')}{A(\sigma')} = \frac{F(\tau')}{A(\tau')}$. Conversely, among the functionals invariant under isometries, only the ones proportional to the area have the limit defined for the Delaunay triangulations of all Delaunay sets.

Subclasses. We are interested in two subclasses of functionals, $\mathcal{G} \subseteq \mathcal{F} \subseteq \mathcal{E}$, which we now introduce. To define \mathcal{F} , let Y be a generic set of $d + 2$ points in \mathbb{R}^d such that

no point lies inside the convex hull of the others. The non-degenerate d -simplices spanned by the points cover the convex hull twice; see Radon [10]. Indeed, we can split them into two collections such that each forms a triangulation of Y : the Delaunay triangulation, $D = \text{Del } Y$, and the other triangulation, T . Changing one triangulation into the other is a *flip*, a name motivated by the planar case in which it replaces one diagonal of a convex quadrilateral with the other. We give the flip a direction, leading from T to D . Let now F be a functional, let Σ_T be the sums of $F(\sigma)$ over all d -simplices in T , and define Σ_D similarly.

DEFINITION. The class \mathcal{F} consists of all functionals $F \in \mathcal{E}$ for which $\Sigma_D \leq \Sigma_T$.

In \mathbb{R}^2 , the extra property of functionals in \mathcal{F} suffices to prove our main result. In \mathbb{R}^d , for $d \geq 3$, we need more structure. The reason is the existence of triangulations that cannot be turned into the Delaunay triangulation by a sequence of directed flips; see [5] for finite examples in \mathbb{R}^3 . Such examples do not exist in \mathbb{R}^2 ; see [7].

Let now Y be a finite set of points in \mathbb{R}^d . As before, we assume that Y is generic. Let T' be a simplicial complex with vertex set Y , but note that we do not require that T' be a triangulation of Y . For example, we could start with a triangulation of Y and construct T' as the subset of d -simplices that do not belong to the Delaunay triangulation together with their faces. Let D' be the subset of simplices in $\text{Del } Y$ contained in the underlying space of T' . Finally, let $\Sigma_{T'}$ be the sum of $F(\sigma)$ over all d -simplices in T' , and define $\Sigma_{D'}$ similarly.

DEFINITION. The class \mathcal{G} consists of all functionals $F \in \mathcal{E}$ for which $\Sigma_{D'} \leq \Sigma_{T'}$.

The condition for F to belong to \mathcal{G} is at least as strong as that for F to belong to \mathcal{F} , which implies $\mathcal{G} \subseteq \mathcal{F}$.

3 Results

In this section, we state and prove our Main Theorem and some of its implications.

3.1 Main Theorem

As already mentioned in Section 1, the main result of this paper is an extension of optimality results for Delaunay triangulations from finite sets to Delaunay sets, which are necessarily infinite. Section 2 provides all the technical concepts needed to give a formal proof of the theorem that facilitates this result.

MAIN THEOREM. Let X be a Delaunay set in \mathbb{R}^d .

- (i) In \mathbb{R}^2 , $F \in \mathcal{F}$ implies $f(\text{Del } X) \leq f(T)$ for all uniformly bounded triangulations T of X .
- (ii) In \mathbb{R}^d , $F \in \mathcal{G}$ implies $f(\text{Del } X) \leq f(T)$ for all uniformly bounded triangulations T of X .

PROOF. Fix a uniformly bounded triangulation T with parameter q of X . We prove the inequalities by comparing subsets of d -simplices of T and $D = \text{Del } X$. We begin with (ii). For every radius α , we write $T(\alpha) \subseteq T$ and $D(\alpha) \subseteq D$ for the sets of simplices contained in \mathbb{B}_α . Furthermore, we write $D'(\alpha) \subseteq D(\alpha)$ for the set of simplices contained in the underlying space of $T(\alpha)$. Summing F over the d -simplices in these sets, we have

$$\Sigma_{T(\alpha)} = [\Sigma_{T(\alpha)} - \Sigma_{D'(\alpha)}] + [\Sigma_{D'(\alpha)} - \Sigma_{D(\alpha)}] + \Sigma_{D(\alpha)}. \quad (14)$$

The first difference on the right-hand side is non-negative by assumption of $F \in \mathcal{G}$. We prove shortly that the second difference is bounded from above by a constant times α^{d-1} . This implies that dividing by $V_d \alpha^d$ and taking the lower limit gives $f(T) \geq f(D)$, as required. To prove the bound for the second term, we assume $\alpha \geq 4q$, so that the Simplex Count Lemma implies that the number of d -simplices in $D(\alpha) - D(\alpha - 2q)$ is bounded from above by a constant times α^{d-1} . We get the same bound for $D(\alpha) - D'(\alpha)$ because $T(\alpha)$ is uniformly bounded, with parameter q , so it covers all of $\mathbb{B}_{\alpha-2q}$, which implies $D(\alpha - 2q) \subseteq D'(\alpha)$.

The proof of (i) is similar, except that we have to do more work to construct the sets of d -simplices, now triangles needed for the comparison. We assume $\alpha \geq 8q$ and let $T(\alpha)$ and $D(\alpha)$ be as before. We construct $T'(\alpha)$ by modifying $T(\alpha)$ through a sequence of directed flips applied to non-locally Delaunay edges. Iterating the directed flip, we can guarantee that all interior edges of $T'(\alpha)$ are locally Delaunay. A directed flip does not increase the size of the larger circumcircle (see e.g. [9]), so flipping does not take us outside the class of uniformly bounded triangulations. Importantly, the flips turn a large portion of $T(\alpha)$ into Delaunay triangles, namely $D(\alpha - 6q) \subseteq T'(\alpha)$. To see this, we note that every triangle $\sigma \in T'(\alpha)$ contained in $B_{\alpha-4q}$ is also in D . Indeed, its circumcircle is contained in $B_{\alpha-2q}$, which is contained in the underlying space of $T'(\alpha)$. If σ were not empty, we would have a vertex inside the circumcircle, which would imply an edge that is not locally Delaunay between this vertex and σ , which is a contradiction; see also [2] where this argument is used to prove part (ii) of the Delaunay Triangulation Theorem. Finally, the triangles of $T'(\alpha)$ contained in $\mathbb{B}_{\alpha-4q}$ cover $\mathbb{B}_{\alpha-6q}$, which implies $D(\alpha - 6q) \subseteq T'(\alpha)$, as required. For the comparison, we consider

$$\Sigma_{T(\alpha)} = [\Sigma_{T(\alpha)} - \Sigma_{T'(\alpha)}] + [\Sigma_{T'(\alpha)} - \Sigma_{D(\alpha-6q)}] \quad (15)$$

$$+ [\Sigma_{D(\alpha-6q)} - \Sigma_{D(\alpha)}] + \Sigma_{D(\alpha)}. \quad (16)$$

The first difference is non-negative, and the second and third differences are bounded from above by a constant times α^{d-1} . Dividing by $V_2 \alpha^2$ and taking the lower limit, as α goes to infinity, we get $f(T) \geq f(D)$, as required.

3.2 Implications in the Plane

There are many functionals on triangles that are known to be in \mathcal{F} . Applying the Main Theorem thus gives many opti-

ality results for Delaunay triangulations of Delaunay sets.

COROLLARY A. Let σ be a triangle in \mathbb{R}^2 , with edges of length a, b, c , let $c_1 > 0$ and $c_2 \geq 1$ be constants, and consider the following list of functionals:

- $F_1(\sigma) = \text{Circumradius}^{c_1}(\sigma)$.
- $F_2(\sigma) = \text{Circumradius}^{c_2}(\sigma) \cdot \text{Area}(\sigma)$.
- $F_3(\sigma) = -\text{Inradius}(\sigma)$.
- $F_4(\sigma) = (a^2 + b^2 + c^2)/\text{Area}(\sigma)$.
- $F_5(\sigma) = (a^2 + b^2 + c^2) \cdot \text{Area}(\sigma)$.
- $F_6(\sigma) = \|\text{Centroid}(\sigma) - \text{Circumcenter}(\sigma)\|^2 \cdot \text{Area}(\sigma)$.

Then $f_i(\text{Del } X) \leq f_i(T)$ for every Delaunay set $X \subseteq \mathbb{R}^2$, for every uniformly bounded triangulation T of X , and for $1 \leq i \leq 6$.

PROOF. It is easy to see that all listed functionals belong to \mathcal{E} . For finite sets, the optimality of the Delaunay triangulation for f_1 and f_4 was proved in [8], for f_2 and f_6 it was proved in [9], for f_3 it was proved in [6], and for f_5 it was proved in [11]. It follows that $F_i \in \mathcal{F}$ for $1 \leq i \leq 6$, so the claim follows from (i) in the Main Theorem.

3.3 Implication in d Dimensions

We have one example of a functional on d -simplices that is in \mathcal{G} , namely the extension of F_5 to three and higher dimensions. Writing a_1 to a_k for the lengths of the $k = \binom{d+1}{2}$ edges of a d -simplex σ , we define $F_R(\sigma) = \text{Vol}(\sigma) \sum_i a_i^2$; see also [1]. Rajan proved that for finite sets in \mathbb{R}^d , the density of F_R attains its minimum for the Delaunay triangulation. We will extend his proof to show that F_R belongs to \mathcal{G} . With this, we get another consequence of the Main Theorem.

COROLLARY B. We have $f_R(\text{Del } X) \leq f_R(T)$ for every Delaunay set $X \subseteq \mathbb{R}^d$ and for every uniformly bounded triangulation T of X .

PROOF. The main tool in this proof is the lifting of a point $y \in \mathbb{R}^d$ to the point $y^+ = (y, \|y\|^2) \in \mathbb{R}^{d+1}$, an idea that goes back to Voronoi [12]. Note that y^+ lies on the graph of the function $\varpi : \mathbb{R}^d \rightarrow \mathbb{R}$ defined by $\varpi(x) = \|x\|^2$. For $Y \subseteq \mathbb{R}^d$, write Y^+ for the corresponding set of lifted points, and let $\text{conv } Y^+$ be its convex hull. Assuming Y is finite and generic, $\text{conv } Y^+$ is a convex polytope whose faces are simplices. We distinguish between *lower* faces whose outward normals point down – against the direction of the $(d+1)$ -st coordinate axis – and *upper* faces whose outward normals point up. Importantly, if we project all lower faces vertically to \mathbb{R}^d , then we obtain the Delaunay triangulation of Y .

Consider now the functional F_L that maps a d -simplex σ in \mathbb{R}^d to the $(d+1)$ -dimensional volume between the convex hull of the $d+1$ lifted vertices and the graph of ϖ . More precisely, it is the volume of the portion of the vertical $(d+1)$ -dimensional prism over σ that is bounded above by the convex hull of the lifted vertices and below by the graph of ϖ . It is not difficult to prove that F_L is invariant under isometries of \mathbb{R}^d , and to use this fact to show that $F_L \in \mathcal{E}$. The reason for our interest in F_E is the relation

$$F_R(\sigma) = (d+1)(d+2)F_E(\sigma) \quad (17)$$

proved in [11]. Since the two functionals differ only by a multiplicative constant, it follows that F_R also belongs to \mathcal{E} . It remains to prove that F_R belongs to \mathcal{G} , which we do by showing that F_E belongs to \mathcal{G} . Indeed, this should be clear from the lifting result: the lifted images of the d -simplices in the Delaunay triangulation are closer to the graph of ϖ than those of other d -simplices. More specifically, if T' is a simplicial complex with finite vertex set Y in \mathbb{R}^d , and all simplices of $D' \subseteq \text{Del } Y$ are contained in the underlying space of T' , then the total $(d+1)$ -dimensional volume we get for T' is larger than or equal to that we get for D' . But this implies $F_E \in \mathcal{G}$, and therefore $F_R \in \mathcal{G}$, as required.

4 Discussion

The main contribution of this paper is an extension of optimality results that hold for Delaunay triangulations of finite sets to Delaunay sets, which are necessarily infinite. In the plane, this extension holds for all functionals that improve upon flipping an edge that is not locally Delaunay. In three and higher dimensions, we need stronger properties to prove the extension. It would be interesting to know whether these stronger properties are necessary. Specifically, is it true that $F \in \mathcal{F}$ implies that the density of F attains its minimum at the Delaunay triangulation of a Delaunay set in \mathbb{R}^d , also for $d \geq 3$? Similarly, are there functionals in \mathcal{G} that are not in \mathcal{F} , or is $\mathcal{F} = \mathcal{G}$?

References

- [1] A. V. AKOPYAN. Extremal properties of Delaunay triangulations. *Trudy ISA RAS* **46** (2009), 174–187.
- [2] B. N. DELAUNAY. Sur la sphère vide. *Izv. Akad. Nauk SSSR, Otdelenie Matematicheskikh i Estestvennykh Nauk* **7** (1934), 793–800.
- [3] B. N. DELONE. Geometry of positive quadratic forms. *Uspekhi Mat. Nauk* **3** (1937) 16–62.
- [4] N. P. DOLBILIN, O. R. MUSIN AND H. EDELSBRUNNER. On the optimality of functionals over triangulations of Delaunay sets *Uspekhi Mat. Nauk* **67** (2012), 189–190.
- [5] B. JOE. Three-dimensional triangulations from local transformations. *SIAM J. Sci. Statist. Comput.* **10** (1989), 718–741.
- [6] T. LAMBERT. The Delaunay triangulation maximizes the mean inradius. In “Proc. 6th Canad. Conf. Comput. Geom., 1994” 201–206.

- [7] C. L. LAWSON. Software for C^1 surface interpolation. In *Mathematical Software III*, Academic Press, New York, 1977, 161–194.
- [8] O. R. MUSIN. Properties of the Delaunay triangulation. In Proc. 13th Ann. Sympos. Comput. Geom., 1997” 424–426.
- [9] O. R. MUSIN. About optimality of Delaunay triangulations. *Geometry, Topology, Algebra and Number Theory, Applications*, Internat. Conf. dedicated to 120th anniversary of B. N. Delone, 2010, 166–167.
- [10] J. RADON. Mengen konvexer Körper, die einen gemeinschaftlichen Punkt enthalten. *Math. Ann.* **83** (1921), 113–115.
- [11] V. T. RAJAN. Optimality of the Delaunay triangulation in \mathbb{R}^d . *Discrete Comput. Geom.* **12** (1994), 189–202.
- [12] G. F. VORONOI. Nouvelles applications des parametres continus a la theorie des formes quadratiques. *J. Reine Angew. Math.* **34** (1908), 198–287.